Performance Benchmarks for Passive UHF RFID Tags*

Karthik Moncombu Ramakrishnan[1] and Daniel D. Deavours[1]

Information and Telecommunications Technology Center
2335 Irving Hill Road
Lawrence, KS
{deavours,karthikr}@ittc.ku.edu

Abstract. This paper describes the first comprehensive benchmark suite for passive UHF RFID tags. These benchmarks give good indications as to how well RFID products will work in "real world" scenarios. We present benchmarks and give some experimental results, as well as key insights that these benchmarks have revealed about the nature of current UHF RFID tags on the market.

1 Introduction

Radio Frequency Identification (RFID) is revolutionizing the way products and goods are tracked and traced in the supply chain. The main reason for that is the potential that RFID offers to businesses in terms of efficiency and new cost-saving capabilities. Retailers consider RFID as an investment for the future, providing advantages like cost reduction by maintaining stock levels, reducing the out-of-stocks, counterfeit protection, "shrinkage" protection, and real-time tracking of supplies.

Recent mandates and recommendations from various retailers (e.g., [1]) and government agencies like Department of Defense (DoD) [2, 3] are requiring their suppliers to use RFID. Also, organizations such as Food and Drug Administration (FDA) are encouraging the pharmaceutical companies to use RFID [4]. This has caused RFID to become important to a large number of people who were unfamiliar with the technology. An estimated 10,000 companies supplying a major retailer and 50,000 suppliers to U.S. DoD have to meet aggressive time lines set by these mandates and recommendations. Few companies affected by these mandates have the necessary in-house RF expertise to deploy the technology. The majority of the companies generally resort to outside expertise for information and help. These companies employ the RFID vendors or third party solution providers to investigate RFID performance in their environment. When companies employ vendors, there is an obvious risk of getting biased information. Even as late as 2004, there were not sufficient third party solution providers having access to good information. Although there are now better third party

* This work was funded in part by RFID Journal and Rush Tracking Systems.

solutions, the risk of bias still exists. Hence, there is a need for unbiased, good, and reliable source of information for RFID products.

The companies that are affected by the mandates need to deploy RFID, which has created new levels of demand for RFID products. The ignorance within the market about the RFID technology and the enthusiasm of RFID vendors has resulted in competing and misleading claims of product performance. For example, one RFID tag vendor currently states on their web page that "Today's RFID tags have read rates varying from as low as 20 tags/second to over 1,000 tags/second." We believe that this statement is false, or at best, misleading. Misleading claims and the lack of good, credible, unbiased source of information has created confusion among the RFID end-users. Hence the need for quality information that can be obtained by quality, fair benchmarks.

RFID Alliance Lab [5] was created to provide unbiased, reliable, and independent source of information for performance of EPC-complaint UHF RFID products (the kind being mandated). The lab provides objective benchmarking information that separates facts from hype. For example, we have observed tag read rates ranging between 0 and 65 tags per second (the details are described in this paper), not the 20 to 1000 as one vendor claims. The RFID market experiencing is not unlike those in the computer market before standardized computer benchmarks were developed [6]. Benchmark information, like tag read rates, will help the end-users to make more informed decisions and deploy RFID technology successfully.

In this paper, we present benchmarks that compare the read performance of different tags in terms of distance, quality, and read rates in various situations. These measures are designed to be both relevant and (to the degree possible) intuitive to the end users. The benchmarks are designed to be scientific and repeatable, as well as indication of the real-world performance of the tags. These measures give information that will aid implementing better RFID systems.

In addition to describing our benchmarks, we give some select results. Space limitations constrain what we can present here; the full results are commercially available [7, 8]. Instead, we focus on tag read performance, and give select data that gives insight into some of the most interesting aspects of RFID tag performance. For example, we observed that the Class 0 tags in large populations can be read considerably faster than Class 1, whereas Class 1 tags when read individually can be read faster than Class 0. This benchmark suite provides a first step towards a common benchmark standard that aims to reduce confusion prevalent among the end users and provide consistent information using user-relevant performance measures.

This paper is organized as follows: Section 2 gives an overview of RFID and its performance, Section 3 discusses about the developed performance metrics for RFID tags in air and near materials, Section 4 briefly describes some other benchmarks and other ongoing work, and Section 5 concludes the work.

2 Background

One of the benefits of an RFID systems is the ability to provide automatic identification to physical objects without the need for line-of-sight communication. The main components of a RFID system are: tags, readers, and host computer. RFID tags are attached to physical objects as the means to identify them. RFID readers convert the radio waves reflected from the tags to get the digital data and send the collected data to the host computer. RFID tags used in supply chain carry a unique serial number called Electronic Product Code (EPC) [9]. Mandates require that the tags deployed in the supply chain to be UHF EPC-compliant tags. In this paper we provide benchmark metrics for passive UHF EPC-compliant tags. (Other non-passive UHF EPC-compliant tags such as those with a battery assist are emerging, but have not yet been studied.)

The reader should note that no so-called "Gen 2" tags [10] are presented in this paper, and it is widely believed that "Gen 2" tags will supplant the current generation of tags shortly. All the benchmarks we discuss here also apply to "Gen 2" UHF RFID tags.

The passive UHF EPC-compliant RFID tags used in the supply chain form a small portion of variety of RFID systems possible. For perspective, we give a brief description of the larger RFID universe before describing passive UHF RFID tags in more detail.

2.1 Taxonomy of RFID Systems

The variety and the operating principles of RFID systems have enabled classification along several dimensions. RFID tags can be classified as chip and chip-less tags, based on whether they include an integrated circuit. The power needed to operate the IC is derived either from the reader's RF signal or from a battery. Chip-less tags do not contain an integrated chip; rather, they encode unique patterns on the surfaces of materials, e.g., Surface Acoustic Wave (SAW) RFID tags which are based upon piezoelectric effects [11].

The Auto-ID center has provided a layered class structure [12] to classify RFID tags based on their operation and functionality. The classification system ranks tags from the least sophisticated (Class 0) to the most sophisticated (Class 5). Class 1 tags represent basic capability. They may be written only once but read multiple times, and obtain all power from the reader's RF transmission. They communicate back to the reader by modulating the chip load impedance, which changes the tag's radar cross signature, a technique called *backscatter modulation*. Later, a Class 0 was added to represent read-only tags, whose identifiers are programmed by the manufacturer. Class 2 tags are passive tags with additional functionality like encryption or memory, and may be writable multiple times. Class 3 tags are semi-passive tags, which means that they have a battery source for operating the IC, but have no transmitter and hence rely on backscatter modulation. Class 4 tags are "active tags," which have a battery source and a transmitter. They may be capable of broadband peer-to-peer com-

munication with other active tags. Class 5 tags are devices that can power other tags as well as communicate with other Class 4 tags, such as readers.

Another major classification dimension is the frequency at which the RFID systems operate. RFID systems generally operate in specific Industrial Scientific Medical (ISM) bands that occupy portions of spectrum from low frequencies like 125 kHz to microwave frequencies like 5.8 GHz and up. The mandates require that the RFID systems be operated in the UHF (Ultra High Frequency) frequencies occupying the ISM bands in 865–960 MHz according to frequency regulations in different countries. The UHF frequency range was chosen for these mandates for maximizing tag read distance.

2.2 Passive RFID System

Recall that passive RFID tags acquire energy from the RF signal of the reader. The basic working of the system can be explained as follows. A passive RFID tag is composed of a chip, an antenna typically on top of a substrate, and may contain a paper label.

The RFID reader sends out RF energy in attempt to read tags. The tag antenna is tuned to receive the RF energy. A voltage multiplier, rectifier, capacitor, and Zener diode are used to convert voltage induced on the antenna to a usable power supply. If enough energy is available, the tag begins to perform the demodulation and processing the commands sent by the reader. The tag responds to the issued commands by switching the load impedance at the antenna terminals from matched to unmatched conditions according to the tag response signal. Changes in antenna load affect the amplitude and phase of the radar cross signature (RCS) of the antenna and are used to encode data. When there are multiple tags responding to a command, the RFID air-interface protocol has an anti-collision algorithm to detect collisions [9, 10, 13].

2.3 Performance Factors

There are two fundamental properties of RFID performance: the fraction of times in which a tag responds, and the speed in which it responds. The former metric is estimated by the ratio of tag responses per requests, and the second is estimated by the number of responses per time. There are a number of factors that impact those values. Examples include the distance between the tag and reader antenna, whether there is metal or other objects obstructing line of sight, whether the tag is placed near a high dielectric material such as water, how much background noise is in the environment, and the orientation of the tag.

Although RFID performance is critically important for successful implementation and there is a lack of quality data available, we are not aware of any published standard or recommendations towards a well-defined set of performance measures. EPCglobal has realized that RFID performance is an issue and is taking steps towards developing a performance standard, but to date there is no publicly available information. The only published previous work we are aware of is [14], which essentially lists a set of simplistic way to compare different

RFID product offerings by end-users. Although it addresses some of the broader performance issues, it is not well developed and is not useful as a benchmark. In this work, we provide well-defined set of measures for performance comparison of RFID tags as well as provide empirical results based on those benchmarks and insights from those results.

While the combinations of factors are numerous, our objective is to define a small set of benchmarks that are meaningful and somewhat orthogonal. A meaningful benchmark is one that gives useful data that is transferable to real-world scenarios so as to give RFID implementors indicators as to how well RFID products will perform. Orthogonal benchmarks are desirable so as to minimize the amount of data necessary for one to comprehend the performance characteristics of a tag. Thus, we took a pragmatic approach, and worked closely with implementors to devise these benchmarks.

3 Benchmark Suite

The time constraints to meet the mandates and the lack of good, reliable, unbiased source of information has driven the need for developing a set of common benchmarks for comparing performance of tags. Many of the performance benchmarks we have developed for comparing tags are presented in this section. A more detailed and formal description of the benchmarks are available [15].

3.1 Response Rate vs. Attenuation

One of the first questions people tend to ask about RFID tags and readers is how far away a tag can be read. This benchmark addresses that question, albeit somewhat indirectly. Figure 1 shows the common setup we have used for all our free space experiments, with d equal to approximately one meter. In order to achieve repeatable results, changing d was simulated by changing the reader's transmitter power levels. Ideally, both the forward channel (reader to tag) and the reverse channel (tag to reader) should be attenuated for simulating distance, but in Section 4.1 show conditions under which attenuating the forward channel is sufficient.

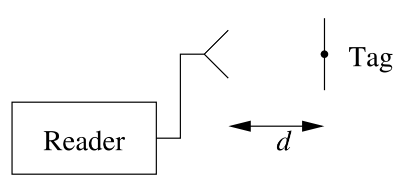

Fig. 1. Free space experimental setup.

The *response rate* is the ratio of the number of successful reads per the number of read attempts. For this test, the tag is kept at a fixed distance from the

reader as shown in Figure 1. The forward channel was attenuated as a means to simulate change of distance. We have translated dB of attenuation into approximate distances using the well-known free-space transmission equation. These results give good, quantitative, relative performance data.

Figure 2 shows an example of a response rate vs. attenuation for a commercial UHF RFID tag.[1] We performed 1 million read attempts over 200 different power settings to obtain this data. We observe three regions of operation: strong-in-field, weak-in-field, and out-of-field. Typically, the response rate is nearly 100% when the tag is in the strong-in-field region, but we have observed response rates as low as 85%. When in the weak-in-field region, the tag exhibits a non-monotonic decrease in response rate. The example in Figure 2 shows a relatively smooth decrease in the response rate, but other tags give a much more "bumpy ride down." In the out-of-field region, the tag response rate is 0%. Early experiments showed response rates slightly larger than 0% for Class 0 tags, which we determined was the result of "ghost tags" (see Section 4.2).

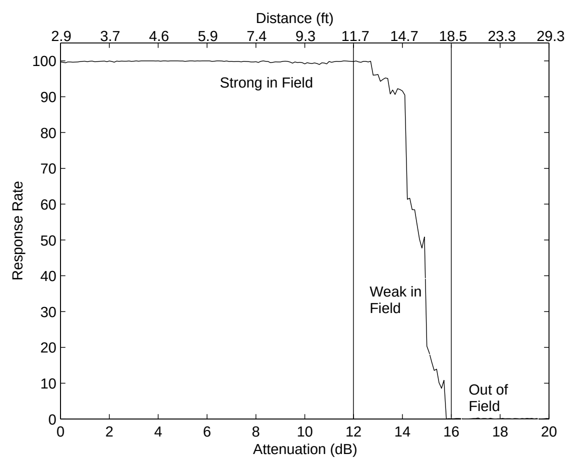

Fig. 2. Typical response rate versus attenuation for a RFID tag.

3.2 Orientation Sensitivity

Clearly, the radiation pattern of the RFID tag antenna is going to affect how well the tag is readable in different orientations. This benchmark was developed to measure that directly. We broadly classify tags into two categories: the "long and thin" tag, and the "squarish" tag. The "long, thin" tags are typically some

[1] Contractual constraints prohibit us from identifying particular RFID products by name, but that information is commercially available in [7, 8]. We believe many lessons and trends can still be learned without identifying particular products.

variant of a dipole, while the "squarish" tags are typically dual dipoles. The radiation pattern of each of these is obviously going to differ. To test the radiation pattern, or orientation sensitivity, of tags, we placed each tag in our free space setup configuration (Figure 1). We rotated each tag at 20° steps along the E- and H-plane, and attenuated the power setting of the reader until the response rate went to 0% in 100 read attempts.

In the H-plane, all but one tag we tested had nearly circular patterns, and thus we present no data. However, there were considerable differences in the E-plane. Figure 3 shows sample data from two commercially available UHF RFID tags. The concentric circles are labeled with dB of attenuation (top) and an approximate conversion to distance in feet (bottom). One can see that the "long and thin" tag exhibits typical "figure-8" dipole behavior, while the "squarish" tag exhibits more uniform radiation. "Long and thin" tags have the advantage of a smaller form factor, while "squarish" tags show an advantage in a near-uniform radiation pattern.

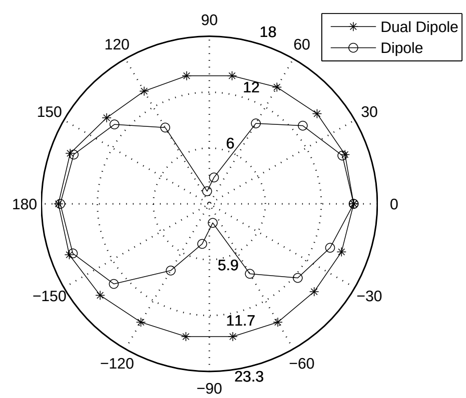

Fig. 3. Orientation sensitivity of two tags along E-plane.

3.3 Variance of Tag Performance

One would normally assume that two tags of the same model would exhibit nearly identical performance behavior. If true, then the RFID system will give predictable performance from tag to tag, and the predictability allows one to engineer high quality processes. Unfortunately, we have found that tag performance varies considerably. We begin by a discussion of how we compare the performance two tags.

We use the response rate against vs. attenuation (see Section 3.1) as the basic performance characteristic for a given tag. One way to compare read performance

is to identify the point in which the read rate goes to 0%, which is what we frequently do. However, we have observed that the 0% point can sometimes yield results that are off by as much as 1 dB. A better way to compare tags using a single number is to use the response curve (e.g., Figure 2) and then compute $\min_\delta ||f_1(x) - f_2(x + \delta)||$, where f_1 and f_2 are the functions representing the response rate vs. attenuation for the two tags, and $||\circ||$ is a norm. We found the 1-norm to be sufficient.

To determine the variance in performance of tags, we measure the response rate vs. attenuation for at least 100 tags of each model. (The number can be adjusted for desired statistical accuracy.) To expedite the testing process, we first chose a small number of tags to determine the range of "interesting" attenuation values. We focused on the range that would include all three regions (strong-in-field, weak-in-field, and out-of-field) for most tags. We then varied the output power in steps of 0.5 dB and performed between 50 and 100 read attempts per power setting (50 attempts were used for "slow" tags for pragmatic reasons).

An example of one commercial tag model that we tested is shown in in Figure 4. We calculated various ranges in tag performances. All the tags tested in each tag model were ranked from 0% corresponding to worst performing tag to 100% corresponding to the best performing tag in that model. The black line depicts the median tag (50%) in this tag model. The darkest band near the median tag shows the middle 40% of tags of the tag model that we had tested. The middle 70% is shown in the next darker color band. The middle 87% and 98% are shown in the next two bands, and the lightest band encompasses all the tags that were tested.

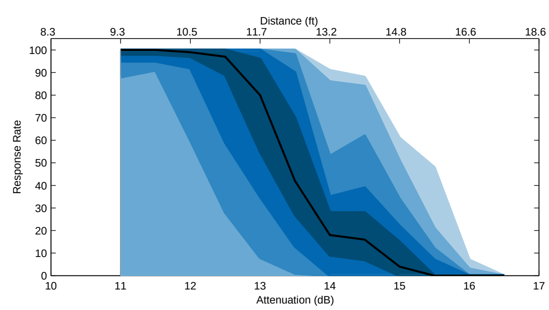

Fig. 4. Variance of performance for a tag model

We note that the tag model depicted in Figure 4 shows considerable variation in performance. Note that the 98% band extends all the way to 11 dB, so the bottom 1% of tags were unreadable at 11 dB of attenuation, while a number of other tags were barely readable at 16 dB. However, the vast majority of tags fall within a 3.5 dB window of performance, or roughly a factor of two. Other tags we tested had different distributions, but similarly large variances. In order

to quantify the variance of performance in a tag model, the norm-metric can be used between the best and worst performing tag on that model.

End users can use this data to compute their own fitness metrics. For example, an end user may look at the fraction of tags (yield) that will give some minimum performance level. Other end users may be concerned about consistent performance to prevent a reader from reading tags going through adjacent conveyors or dock doors, for example.

3.4 Read Rates

Recall our example from Section 1 of how a vendor claims that tags are readable from between 20 and 1000 times per second, a claim we believe is false or a best misleading. This benchmark is designed to put to rest some of this marketing "hype." *Read rate* is defined as the ratio of the number of times a tag is read per number of seconds reads were performed. This is one of the two fundamental characteristic of the RFID tag-reader system, and denotes the speed at which the reader can read a tag. The benchmarks that are based on this characteristic are: read rate in isolation, tag read rate in population, time to first read, and total read rate.

Read Rate in Isolation When only a single tag can be detected by the reader, we say the read rate is performed *in isolation.* We chose a median-performing tag (from the study in Section 3.3) for each tag model. The reader was programmed to read tags as many times as possible for 60 seconds. The reader setting was set to look only for Class 0 or Class 1 tags accordingly, and the timeout value was determined using the reader guidelines. The experiment was repeated 10 times for each model.

One interesting finding is that we found two types of Class 1 tags, "fast" and "slow." Within the two types, tag models had similar read rates, but the difference between the two were remarkable (see Table 1). We can speculate that different Class 1 tag models use different integrated circuits. Class 0 tags were comparable to "slow" Class 1 tags. The standard deviation in read rates were modest, indicating the read rates are robust.

Table 1. Read rates in isolation.

Tag class	Tag type	Rate	St. Dev.
Class 1	"Slow"	7.0	0.24
Class 1	"Fast"	24.1	1.29
Class 0	-	6.5	0.04

We repeated this experiment with reader from other manufacturers, and for brevity we summarize our observations we found. Different readers showed similar trends. However, the absolute values of read rates can vary by as much as 250%. Thus, the read rates are very much dependent on the tag-reader system.

We note that some readers implement a feature of the EPC Class 1 protocol called "ScrollAllID," or simply "global scroll." This feature allows the reader to bypass any collision detection and resolution algorithms, and can consequently yield read rates of over 400 tags/sec. Since the global scroll mode avoids collision detection, we generally do not recommend its use except in very well-controlled conditions, and hence we do not use it for this benchmark.

3.5 Read Rate in Population

In general, there can be multiple tags present in the reader's read field. Typical applications for this include reading all the items within a container, and reading all the containers on a pallet. Depending on the air-interface protocol between the tag and the reader, there are different approaches to collision detection and resolution. Class 0 uses a binary tree search approach [13]. This means that the EPC is represented as a binary tree and the reader scans the binary tree from the root to leaf in order to identify a tag. This process of identification of a tag is called tag singulation. When the reader is reading Class 1 tags, the reader puts the tags that are already read to "sleep" so that the reader can focus on reading the other difficult-to-read tags [9]. Typically, a RFID reader performs a sequence of "wake up", "read", and "sleep" cycles to ensure that all Class 1 tags in the field are read. It will become apparent that these different approaches in handling collision detection and resolution can result in an order-of-magnitude performance difference between tags.

The fundamental aim of this test is to populate the read field with as many tags as possible. That turned out to be challenging because tags placed too close to each other can degrade performance considerably, and tags can also mask other tags behind them. Since the tag size and interaction was unique to each tag model, the placement in space is unique, and thus the results are not completely reproducible by others or comparable across models. Also, the goal of putting many tags in the read field means that many or most may be in the weak-in-field region of operation.

Despite the drawbacks, this test shows very different trends in the Class 0 and Class 1 approach to populations of tags. We performed this experiment with three tag models: a Class 1 tag with 140 tags in the read field, a Class 0 tag with about 120 tags, and 48 Class 0 item-level (small) tag on a bottles of placebo stacked in a cardboard box. We performed the test 10 times on each of the three tags. We obtain three metrics from this test: time to first read (TTFR), total read rate in population, and individual read rate in population. Each of those metrics are defined below.

Time to First Read The *time to first read* (TTFR) metric is defined as the time it takes to read the n-th tag for the first time, which can vary considerably. Figure 5 illustrates the average TTFR for the Class 0 and Class 1 tag. (The Class 0 item tag was not included in this analysis.) We limited this test to 300 seconds (5 minutes), and if a tag was not read after 300 seconds, we assigned a

value of 1,000 seconds (so it could be plotted). We drew a horizontal line at 300 seconds in Figure 5 for convenience. It should be noted that all tags were read at least once in at least one repetition.

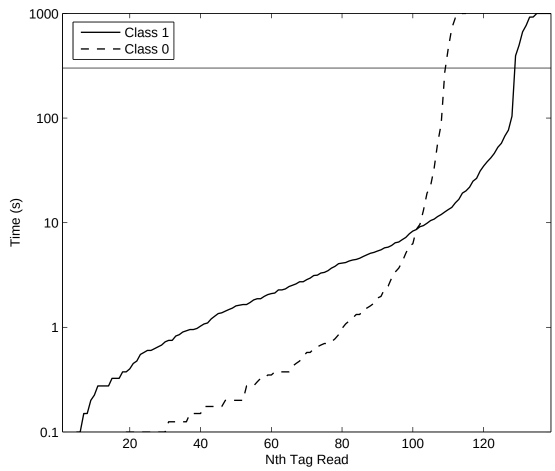

Fig. 5. Time to first read (TTFR).

First, we note that there was variance in the TTFR among the different repetitions, but was not plotted to keep the figure clear. The variance was small and does not affect the conclusions we draw from the results. We see that for the first two-thirds of tags, the TTFR for Class 0 is significantly smaller than the TTFR for Class 1. This matches the general observation that Class 0 performance in population is much better than Class 1. For larger n, Class 0 was slower, but this is likely because there were fewer Class 0 tags tested, or because of other experimental artifacts.

The most interesting observation is that both tags exhibited the same general trend: the first few tags were read rapidly, about two-thirds of the tags were read in nearly linear time, and the last third of tags shows an almost exponential growth in the TTFR. We lack the instrumentation to understand why this is the case, and note it as an interesting topic for future work, especially since "Gen 2" tags are based on the Class 1 protocol.

Total Read Rate in Population The *total tags read rate in population* is defined as the number of times any tag in the population was read divided by time (300 seconds). The results for the three tags tested with this benchmark are given in Figure 2. Notably, Class 0 tags yields a total read rate in population of between 4.7 and 5.8 times greater than Class 1 under this benchmark.

Table 2. Total Read Rate in Population.

Tag	Rate	St. Dev
Class 1	45.6	0.99
Class 0	265.5	8.77
Class 0 item	212.3	4.80

Individual Read Rate in Population In a population of tags, some tags are read more frequently than others. We define the *individual read rate in population* of a tag is the number of times the individual tag is read per interval of time. Figure 6 illustrates the results when the tags are sorted in decreasing read rates in population.

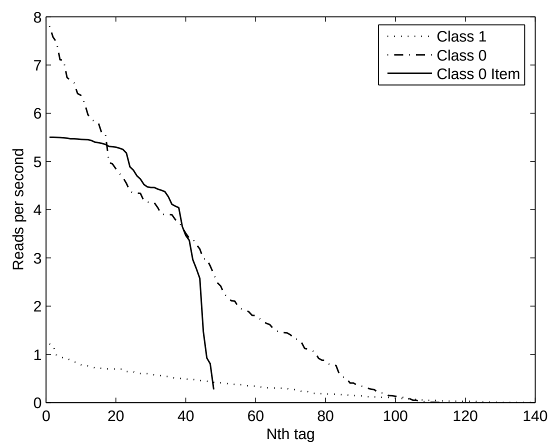

Fig. 6. Tag read rate in population

It is interesting to note that the Class 1 and non-item Class 0 tag follow a similar trend, with Class 0 being proportionally faster. However, the Class 0 item tag shows a very different shaped curve. We conjecture that the reason is that the setup for the Class 0 item tag was such that the majority of tags were in the strong-in-field region. It is also interesting to compare Figure 6 and Table 2 to Table 1. We see that the read rate for a some Class 0 tags in population do not substantially decrease from in isolation, while the Class 1 tags decrease from about 24 to 1.

3.6 Read Performance near Metal and Water

So far, the benchmarks we have discussed have all been in "free air." Free air metrics can be used as baseline measurements for tags, but tags are meant to

be placed on items, and people need to know how tags will respond in those conditions (or whether they will respond at all). The two most common and difficult materials are water and metal. In fact, we show in Section 3.7 that one of the more difficult "real world" products to tag was dishwasher detergent because the product was contained in foil-lined boxes.

Water and metal can affect tag performance in a number of ways. First, they can provide a multi-path, creating fade zones. (They can also be used for constructive interference to boost performance.) Second, the presence of a high dielectric material can change the resonant frequency of the antenna and "de-tune" the antenna. Metals have a similar affect with different mechanisms, reducing antenna efficiency and power transfer efficiency to the chip.

We have developed essentially two variants of a test case involving the presence of materials near tags. The same reader-tag setup of Figure 1 is used, except that planar shaped material is placed some distance l from the tag, and that distance is varied. First, we attenuate the power setting until the reader no longer reads after some number of read attempts or some period of time. The data presented in Figure 7 was obtained for the largest power setting yielding no reads after 100 unsuccessful read attempts for four Class 0 tags, tested for l ranging from 0 to 2 cm in steps of 2.5 mm.

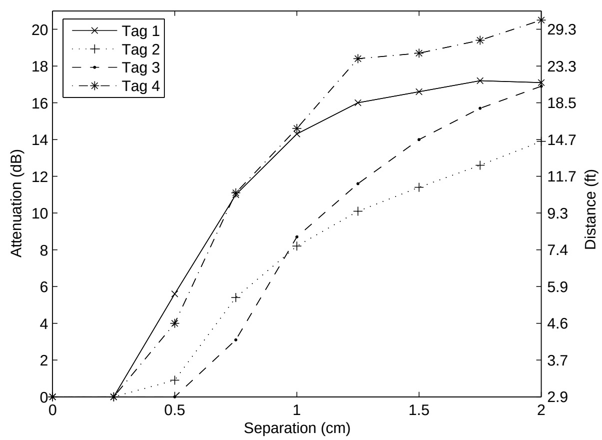

Fig. 7. Tag performance in front of metal.

As Figure 7 shows, there are some significant differences in performance among different tag models, both in how they perform "close" to metal (about 5 mm) and "far" from metal (about 2 cm).

Next, we took essentially the same experimental setup and programmed the reader to perform reads at a single, fixed frequency, and varied the frequency. Note that the ISM band in UHF frequencies varies among countries. The ISM

band frequencies are 865–868 MHz in Europe, 902–928 MHz in USA, 950–956 MHz in parts of Asia. Thus, if a tag is to be read globally, they should be able to perform well across the spectrum. For this benchmark, we performed the experiment at 902, 915, 928, and 955 MHz.

Figures 8a and 8b show two of the most interesting results. We label these two tags "Tag 1" and "Tag 2" respectively. Note that Tag 1 performs better at lower frequencies, and moderately degrades with increasing frequency. This type of frequency variation in performance is typical. However, Tag 2 shows extreme behavior. There are drastic changes in performance with increasing frequency; with 2 cm of separation, there is 16 dB of link margin at 902 MHz (better than Tag 1), but Tag 2 is unreadable at 955 MHz! (865 MHz was not tested due to limitations in our capabilities, but is planned future work.)

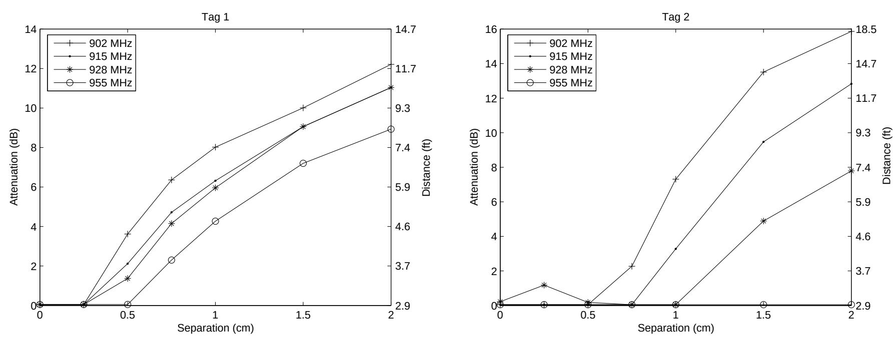

(a) (b)

Fig. 8. Frequency-dependent behavior of Tag 1 (a) and 2 (b).

Future work includes the investigating the rotation of the tag and material to create a radiation pattern, similar to that of Section 3.2, as well as including frequencies between 865–868 MHz.

3.7 Conveyor Testing

The performance metrics that we have discussed until now are based on a controlled, laboratory environment. Laboratory testing is important for scientific, repeatable results, but RFID is not just going to be used in laboratory environments. The greatest current use of UHF RFID is at distribution centers and warehouses to identify palettes going through dock doors and cases traveling on a conveyor. Thus, we also developed a "real world" conveyor benchmark. We note that part of the RFID mandates include the ability to read products passing through a portal with 6 inch spacing between products and the conveyor moving up to 540 to 600 feet per minute. We did not have access to a high-speed conveyor, so we constructed our tests based on available equipment.

Unfortunately, this test is less repeatable than others because there are too many variables that are difficult to constrain. For example, products may shift in cases over time, or differ slightly from case to case, and such small changes can cause significant changes in performance. Further, the number of repetitions necessary for statistical confidence are cost prohibitive. Conveyors come in different dimensions, different sized rollers or sleds, guard rails, and other metal in the vicinity. Other RFID readers may be operating in the area. Thus, no two conveyor setups are identical. Finally, operating a high-speed conveyor with a closed loop can be prohibitively expensive. However, we expect gross features to be useful, and as we show, there are lessons to be learned from this activity.

The experimental begins by placing tags on cases of product. We moved tags around the different faces of the product until we found a position that maximized the read distance. Next, we set up a portal across a conveyor section using three reader antennas (two on the sides, one above). We placed product on the conveyor, spaced about 2 feet apart, and the conveyor was moving approximately 200 feet per minute (considerably slower than RFID mandates). We used between three and five cases of product, adjusted the power level, and performed at least ten reads per power level and product type.

The metric we show here is the the *product read rate*, the fraction of times a product was successfully read when passing through a portal. Figure 9 shows the performance of two different commercially available tags on various products that were tested. The vertical axis represents how much we could attenuate the reader power before the read rate dropped below 50%. (Of course, the 100% value is more interesting, but more difficult to determine with few samples.) Tag 1 is a dipole/slot design, while Tag 2 is a dual-dipole design. Note that Tag 1 was virtually unreadable when placed on dishwasher detergent (dishwasher detergent is in foil-lined boxes), but Tag 2 was unreadable on soup.

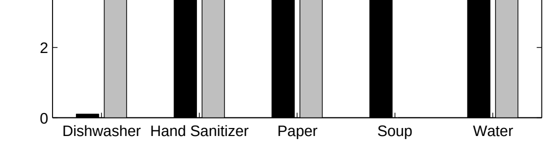

Fig. 9. Performance of two tags in conveyor.

The other interesting item was canned soup. Naturally, a case of canned soup is filled with metal cylinders, with substantial air gaps between the cylinders. It appears that dipoles (”long and thin”) perform well on canned soup because they can be placed on the air gaps. However, one dual dipole design (Tag 2 of Figure 9) did not perform well. We conjecture that the reason for that is that the geometry of the dual dipole is so large that there is always some portion of both dipoles that are being “grounded” by the soup. Even though Tag 2 performed well on the metal test, the interaction of the tag geometry and the metal cylinders afforded no good tag placement.

We note that this benchmark is not well developed, but is of great interest to the RFID community, and thus is a rich topic for future work.

4 Other Benchmarks

Due to space considerations, we only highlight a two other tests that we have performed that we believe are of interest. These include channel sensitivity and ghost tags.

4.1 Channel Sensitivity

For this test, we wanted to see how much the system performance was dependent on the forward channel (reader-to-tag) versus the reverse channel (tag-to-reader). The results of this will illustrate where the constraints are in the current passive UHF RFID systems, and whether attenuating only the transmit power is a good approximation to distance.

We began by placing a tag and reader in the standard setup (see Figure 1). We then placed variable attenuators in the cabling between the reader and both the transmit and receive antenna, and varied both independently by steps of 0.5 dB. The transmit line was attenuated between 12 and 20 dB (no change was observed less than 12 dB), and the receive line was attenuated between 0 and 20 dB. At each setting, we performed 100 read attempts and recorded the response rate. The results of this are plotted in Figure 10. The axis “Tx Att” represents the attenuation in the forward channel, and “Rx Att” represents the attenuation in the reverse channel. The figure shows the results for a Class 1 tag, but we have found nearly identical results for a Class 0 tag.

Although there is some statistical “noise,” it is clear that the response rate showed no sensitivity to attenuation in the receive line. We can infer from this that the tag-reader system is dominated by the forward channel power, and in particular, the ability to transfer sufficient power to the IC. Thus, attenuating the forward channel is sufficient for simulating an increase in distance between the tag and reader. It is interesting that item-level tags behave differently (space constraints limit exploring this further here).

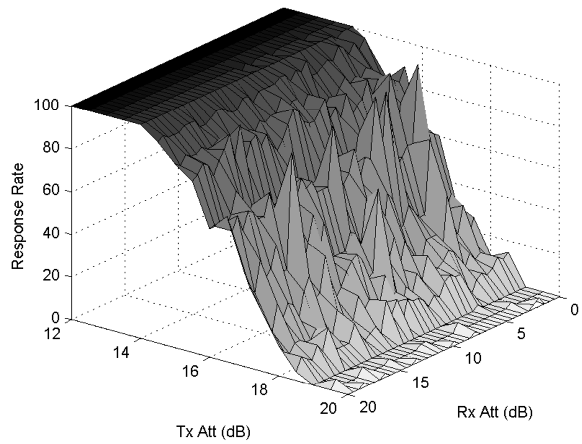

Fig. 10. Effects of forward and reverse attenuation.

4.2 Ghost Tags

When we placed the reader in a mode that would search for Class 0 tags, we noticed from time to time that the reader would identify a tag that was not in the read field. We eventually learned that the ID that was returned was random. It turns out that this is a common problem, so we developed a simple experiment to look for these "ghost tags." In this experiment, we configured the reader for look for any Class 0 tag for the time it would take to perform 100,000 Class 0 reads, and we recorded all the tags that were read and the number of time each tag was read.

We observed 90 distinct Class 0 RFID tags were read by the reader, or 0.9 per 1000. No ghost tag was read more than once. The tag IDs appeared random, and they were evenly split between 64-bit and 96-bit IDs. We also mined data from other experiments from 515,000 Class 0 read attempts and estimate 677 were "ghost tags," or roughly 1.3 per 1000.

Apparently readers using the Class 0 protocol are very susceptible to interpreting noise as signal (due to the FM modulation used). Every tag read is accompanied by a 16 bit CRC, meaning that a random ID and CRC only has 1 in 2^{16} chance of being valid. If 1 out of every 1,000 read attempts, or about 1 in $2^1$0, returns a "ghost tag," then for every read attempt, approximately 2^6 or 64 tags are read with invalid CRCs and are discarded. We find that rate alarmingly high. We have not observed ghost tags with Class 1 tags after looking at data from more than 1 million read attempts. Note that "Gen 2" tags are based on the Class 1 protocol, and thus is not likely to suffer from ghost tags.

5 Conclusion

In this paper, we present a number of simple benchmarks to evaluate the performance of passive UHF RFID tags. We describe each of the benchmarks, illustrate data obtained from performing those benchmarks, and discuss the lessons learned from those results. These results show that there is a variety of different performance between tags, and illustrate the substantial difference between Class 0 and Class 1 tags. This data will enable end users to make informed decisions regarding RFID products, and sort out fact from marketing fiction. We believe that the results from our testing give a better idea to the end-users about which tags would meet their individual needs and help them in implementing better RFID systems.

References

1. Wal-Mart, Inc.: Wal-mart stores — radio frequency identification. http://walmartstores.com/GlobalWMStoresWeb/navigate.do?catg=339 (2006)
2. Office of the Deputy Under Secretary of Defense (Logistics & Material Readiness): Logistics and materiel readiness, home page. http://www.acq.osd.mil/log/rfid/index.htm (2006)
3. Wynne, M.W.: Radio frequency identification (RFID) policy. Policy statement, The Under Secretary of Defense (2004)
4. US Food and Drug Administration: Radiofrequency identification feasibility studies and pilot programs for drugs. http://www.fda.gov/oc/initiatives/counterfeit/rfid_cpg.html (2004)
5. Deavours, D.: RFID Alliance Lab. http://www.rfidalliancelab.org (2006)
6. Corporation, S.P.E.: Standard performance evaluation corporation. http://www.spec.org (2005)
7. Deavours, D.D.: A performance analysis of commercially available UHF RFID tags based on EPCglobal's class 0 and class 1 specifications. Report 1, RFID Alliance Lab, Lawrence, KS (2004)
8. Deavours, D.D.: UHF EPC tag performance evaluation. Report 2, RFID Alliance Lab, Lawrence, KS (2005)
9. EPCglobal Inc.: EPCTM generation 1 tag data standards, version 1.1 revision 1.27. Technical report, EPCglobal Inc. (2005)
10. EPCglobal Inc.: EPC radio frequency identification protocols class-1 generation-2 UHF RFID protocols for communication at 860–960 MHz, version 1.09. Technical report, EPCglobal Inc. (2005)
11. Finkenzeller, K.: RFID Handbook. 2 edn. Wiley & Sons (2003)
12. Sarma, S., Engels, D.W.: On the future of RFID tags and protocols. White paper, Auto-ID Center, Massachusetts Institute of Technology (2003)
13. Auto-ID Center: Draft proposal specification for a 900 mhz class 0 radio frequency identification tag. Technical report, Auto-ID Center (2003)
14. Eberhardt, N.: Towards RFID peformance benchmark tests. Technical report, Auto-ID Center, Massachusetts Institute of Technology (2002)
15. Ramakrishnan, K.N.M.: Performance benchmarks for passive UHF RFID tags. Master's thesis, University of Kansas, Lawrence, KS (2005)